CHAPTER

8

SECURING ACCESS FOR REMOTE ACCESS USERS

After reading this chapter and completing the exercises you will be able to:

- Implement remote access for dial-up clients.
- Implement remote access for VPN clients.
- Implement enhanced security for Windows 2000 remote access.
- Plan and implement remote access policies.
- Understand how Remote Authentication Dial-in User Service (RADIUS) can be used to centralize the management of remote access.
- Implement a corporate remote access plan using Internet Authentication Server.

One of the requirements in most corporations is to give users remote access to the corporate network. In some cases, providing this remote access may be quite simple. For example, a company may decide to give remote access to a single network resource, like e-mail, to only a few executives. The company may limit the remote access even more by giving the executives access to the corporate network only from their home computer and only when they dial into the company's **remote access server (RAS)**. In other cases, the requirement to provide remote access to the network can be much more complicated. A company may have a large number of users who require remote access because they spend most of their time outside the office. For example, a company may have a mobile sales force or executives who travel a great deal. These people require access to the corporate network from almost anywhere in the world. These users may be using a laptop or sitting at a kiosk computer in an airport. In addition to almost universal access, the users may require access to a variety of resources, including e-mail, an intranet Web server, a database server to run a client application, and files on the corporate file servers.

To be complete, your security plan must include policies on how to make the remote access as secure as possible. For most companies, remote access is the only place, other than the Internet firewall, where the network is directly accessible to the outside world. To give remote users access to your network, you have to open an access point through a dial-up connection or through a VPN connection. In either case, you are giving a user direct access to your network, which is exactly what an attacker would like to have. So an essential part of your security plan will deal with making these access points as secure as possible.

This chapter details the requirements for secure remote access. First, this chapter will cover some general remote access concepts and procedures. Then, the remote access authentication protocols and account lockout configuration will be discussed. One of the new features in Windows 2000 is the option to use remote access policies to configure complex conditions that can be used to control when and how remote access clients can connect to the network. Next, remote access policies will be discussed. And finally, this chapter will examine the implementation of Internet Authentication Server, which is Microsoft's implementation of the RADIUS open standard to centralize the administration of remote access for large multinational corporations.

Implementing and Configuring Routing and Remote Access

Before discussing the security configuration for dial-up and VPN remote access, you need to understand the procedures and options for configuring remote access servers and clients. The next sections detail the configuration options for dial-up and VPN servers, as well as the remote access client configuration.

Configuring a Dial-up Server

Many organizations still provide dial-up remote access for users who travel extensively. Despite the fact that dial-up access is slow (with a 56 Kbps maximum connection speed) and often unreliable, dial-up access is often the only option that the remote user has available. Few hotel rooms have high speed Internet connections that you can use to configure a VPN connection to your network; almost all hotel rooms have a data line that you can use to dial out using your laptop's modem. The primary advantage of dial-up access is that it is available almost anywhere.

The first step in providing remote access to clients is to configure the physical modem that the clients will be connecting to. If the users are dialing in using analog modems, you will need to install and configure the modem or modems on the server. If the connection will be an ISDN or X.25 connection, then you will have to configure the hardware device for the connection. In most cases, the modem is detected by plug and play in Windows 2000 and automatically installed during the installation process. However, the modem may not be detected until you scan for hardware changes in Device Manager. If the modem is not detected, then use the Add New Hardware Wizard to add the correct modem. After you have added the modem(s), the next step is to configure Windows 2000 server as a dial-up server. The entire dial-up server

configuration is done through the Routing and Remote Access administration tool. To configure a Windows 2000 server to act as a dial-up server, use the following procedure:

1. Click **Routing and Remote Access** from the Administrative Tools folder.
2. If you have not configured RRAS before, right-click the server name and choose **Configure and Enable Routing and Remote Access**.
3. The Setup Wizard starts. Click **Next**.
4. The Routing and Remote Access Service is used for all remote connectivity options in Windows 2000. The second screen of the RRAS Setup Wizard, shown in Figure 8-1, displays all of the components that you can configure. To configure the server to accept dial-up clients, select **Remote access server**. Click **Next**.

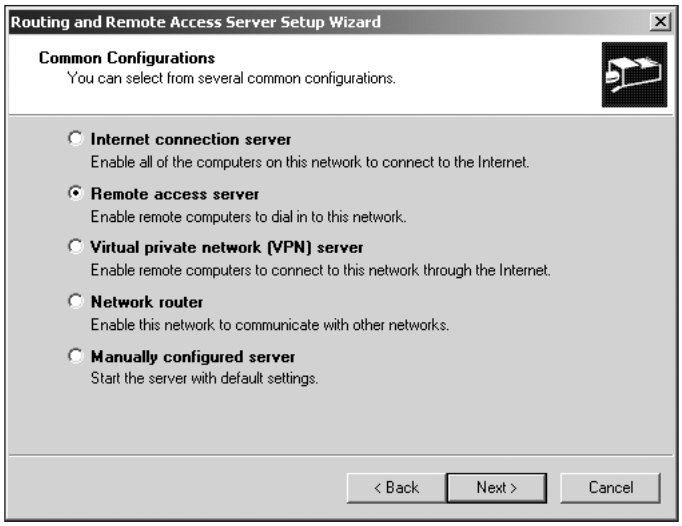

Figure 8-1 Routing and Remote Access configuration options

5. You are then prompted for which protocols the dial-up clients will be using to connect to the server. The protocols must be installed and configured before setting up the remote access server. If all the needed protocols are already installed on the server, click **Next**.

If you have the Apple Talk protocol installed on the server, a dialog box will appear after Step 5 that will enable you to allow unauthenticated access for all remote access clients. This is needed so that Macintosh clients can access the RRAS server. Be sure this coincides with your security policy before enabling this feature.

6. Select how the server will assign IP addresses to the dial-up clients. See Figure 8-2.

Figure 8-2 Configuring RRAS to assign IP addresses to client computers

a. If you are using a DHCP server, then select **Automatically**. The remote access server will then request 10 IP addresses from the DHCP server to assign to dial-up clients. As more addresses are needed, the remote access server will continue to request IP addresses in 10-address increments.

b. If you are not using a DHCP server, you can still select **Automatically** and the server will assign addresses from the Automatic Private IP Addressing range (169.254.0.1–169.254.255.254).

c. If you want to statically assign IP addresses, then select **From a specified range of addresses**. Click **Next**. In the Address Range Assignment dialog box, click **New**, fill in the start and end IP address range, and then click **OK**.

d. Select how you want to assign the IP addresses. Click **Next**.

7. You are prompted to set up the computer to use RADIUS for authentication. If you choose to use RADIUS, then you will be prompted for the name of the RADIUS server and shared secret password. Click **Next**. Click **Finish**.

The RRAS service starts on the server, and then the server is ready to accept dial-up clients. If you have previously configured RRAS on the server and just need to add the option for dial-up clients, right-click the server name and click **Properties**. Ensure that **Remote access server** is selected. Click **OK** and then click **Yes** to restart the RRAS service.

The previous procedure installs the remote access service with a default configuration. There are a number of options that can be configured after installation.

Configuring IP Addressing Options for the Dial-up Server

Depending on the choice that you made during the installation, the server may or may not be using DHCP to assign IP address information to the dial-up clients. If you want to change this configuration after the installation, you can do so by accessing the server properties in RRAS and then selecting the IP tab. See Figure 8-3.

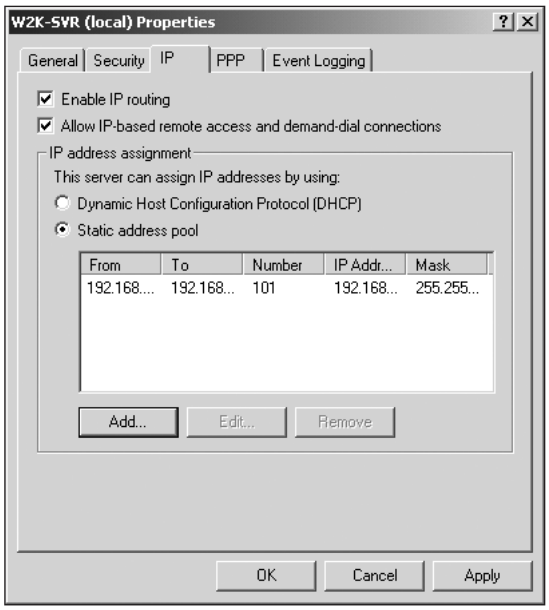

Figure 8-3 Configuring the IP addressing in RRAS

If you have enabled the server to use DHCP and want to change the setting, select Static address pool. You will then need to configure the IP addresses that the server will hand out to dial-up clients. To do so, click Add. Type the IP address range that you want to assign to the clients and click OK. If you are using DHCP to assign IP addresses to the dial-up clients, the clients will receive an IP address and the subnet mask from the RAS server, but not the other DHCP settings, such as WINS server or DNS server IP addresses. In most cases, you should configure the clients to get this additional information. To do so, you install DHCP Relay Agent as a routing protocol, and then add the dial-up interface to the list of interfaces that will act as a DHCP Relay Agent.

Multilink and Bandwidth Allocation Protocol Options

If your server has more than one physical connection that can be used for dial-up access, the Windows 2000 Server can be configured to support **multilinking**, which combines

multiple physical connections into one connection in order to increase bandwidth. You can configure multilinking using multiple modems, ISDN, or X.25 connections. By using the Point-to-Point (PPP) multilink protocol, the client can dial in to the server on multiple lines, and these lines will be combined together into one logical connection.

Windows 2000 Server also uses **Bandwidth Allocation Protocol (BAP)** to help dynamically manage multiple links. For example, if a user is connected on both of the B channels for an ISDN connection, BAP can be configured so that when the use of one of the channels falls below a set limit for a period of time, one of the channels will be dropped, and would then be available to another client.

Both multilink and BAP are configured on the PPP tab on the server properties as shown in Figure 8-4.

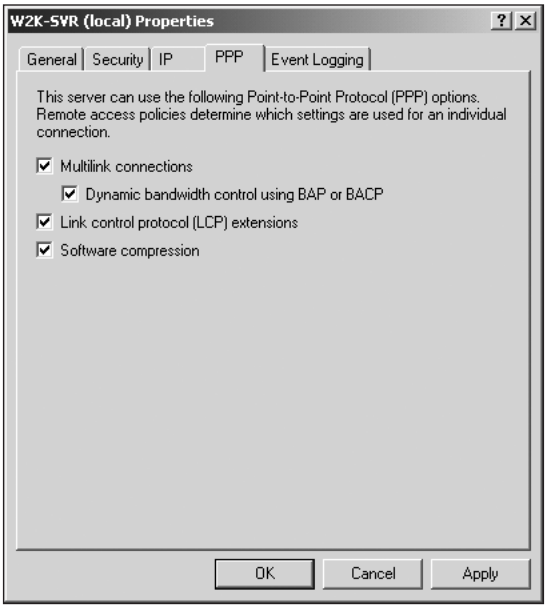

Figure 8-4 Configuring multilink and BAP on the PPP tab for server properties

To configure multilink and BAP, select both options: Multilink connections and Dynamic bandwidth control using BAP or BACP. This is all that is needed to enable the server to use both multilinking and BAP. You will still need to use the remote access policies to further configure multilink and BAP, including setting controls on how many links a client can use and configuring BAP settings.

Configuring Remote Access Clients

After the server has been prepared, the next step is to configure the dial-up clients. The client configuration can be completed either manually or through a Connection

Manager (CM) application. In Windows 2000, individual connections can be manually configured using the Network Connection Wizard.

To configure a dial-up connection manually, use the following procedure:

1. Choose **Start/Settings/Network and Dial-up Connections**.
2. Double-click **Make New Connection**.
3. The Network Connection Wizard starts. Click **Next**.
4. You are given the option of creating several different types of network connections. See Figure 8-5. To configure a dial-up connection, click **Dial-up to private network**. Click **Next**.

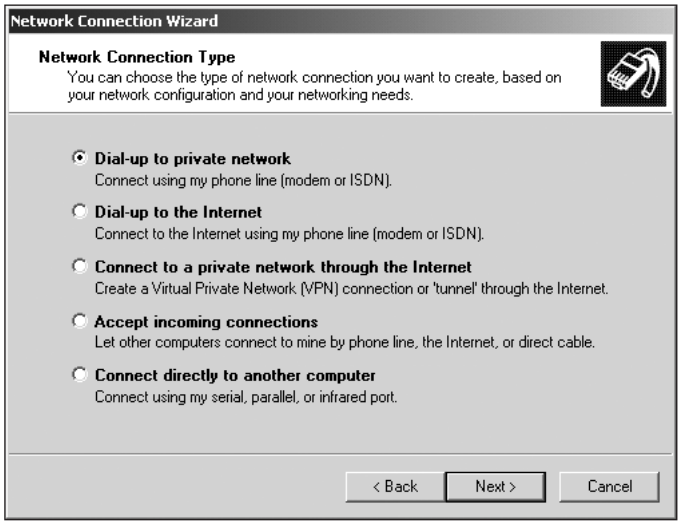

Figure 8-5 Configuring the network connection type for a dial-up client

5. Type a Phone Number, Area Code, and Country/Region Code, if necessary. Click **Next**.
6. Specify whether the connection is for all users or only for a specific user. If you want everyone who logs into the computer to be able to use the connection, then accept the default of creating the connection for all users. Click **Next**.
7. Choose whether you want to enable Internet Connection Sharing (ICS) for this connection. ICS makes it possible for multiple users in a small office to share a single connection to the Internet or another network.
8. Type a name for the connection and click **Finish**.

This procedure installs the network connection with a standard configuration. To manually configure additional settings on the dial-up networking connection, use the following procedure:

1. Choose **Start/Settings/Network and Dial-up Connections**.
2. Right-click the connection you want to configure and click **Properties**.
3. The General tab is used to configure the phone number to be dialed, as well as more advanced settings available, by selecting the following options:
   a. *Configure*—Used to access and configure the modem properties.
   b. *Alternates*—Used to configure additional phone numbers to be dialed if the connection cannot be established on the default phone number.
   c. *Rules*—Used to configure dialing locations. Dialing locations can be used to configure special settings for the dial-up, depending on where you are calling from, including options such as numbers that must be dialed to access an outside line, rules for specific actions when dialing a particular area code, and calling card numbers. To configure locations, click Rules, and then you can add a new location or edit an existing one.
4. The Options tab is used to configure dialing options, including what is presented to the user during the dialing process, as well as how to configure redialing if a dial attempt fails. See Figure 8-6. If you are using this connection with a phone line, you should change the default for Idle time before hanging up. The default of never means that you have to manually disconnect the connection.
5. The Networking tab is used to configure the type of dial-up server you are connecting to, Point-to-Point Protocol (PPP) or Serial Line Internet Protocol (SLIP), as well as the networking components that are used on the connection. See Figure 8-7. The default configuration for the networking components is to use only TCP/IP (configured to obtain an IP address automatically) and activate only the Client for Microsoft Networks. If you are accessing a Microsoft network and do not want to use this computer to share resources, this is the best option. If you want to share files on this computer, you will also have to enable File and Printer Sharing for Microsoft Networks.
6. Click **OK** to accept any changes for advanced configuration.

Implementing Virtual Private Network Access

An increasingly popular alternative to providing the traditional dial-up remote access is providing **Virtual Private Network (VPN)** access. A VPN is used to ensure that your data communication on a public network, such as the Internet, is secure. This is implemented by having the VPN create an encrypted tunnel through the Internet. To create this tunnel, a client computer establishes a PPP connection to an ISP that gives the client access to the Internet. Once the PPP link is established, the client negotiates

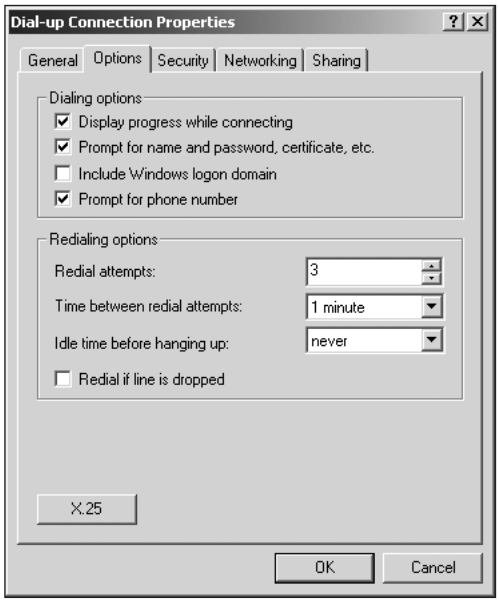

Figure 8-6 Configuring the dial-up options

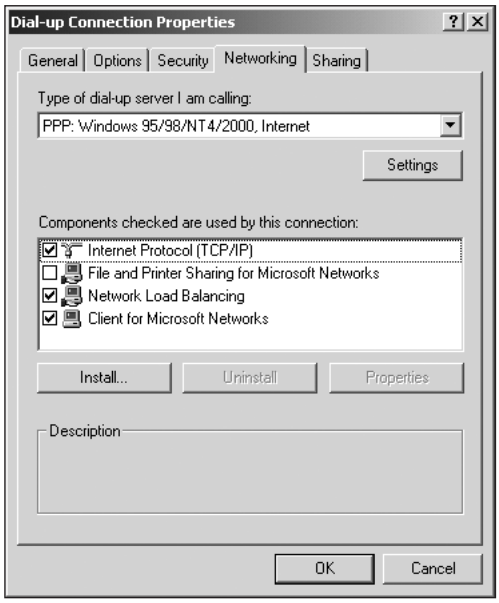

Figure 8-7 Configuring the network options for a dial-up client

a second connection to a remote access server at the corporate network that is also connected to the Internet. In this second negotiation, the client and the server agree on how the packets over the virtual connection will be encapsulated and encrypted. Then the

two computers can send data across the Internet connection with the assurance that the data is secure. The VPN can be created over any existing connection to the Internet. For example, in many parts of North America, cable modem or Digital Subscriber Line (DSL) access to the Internet is widely available and very fast. The VPN can be created using this fast connection to the Internet, providing the client with secure and fast access to the corporate network. Figure 8-8 illustrates how the VPN is created using a dial-up or cable modem connection to the Internet.

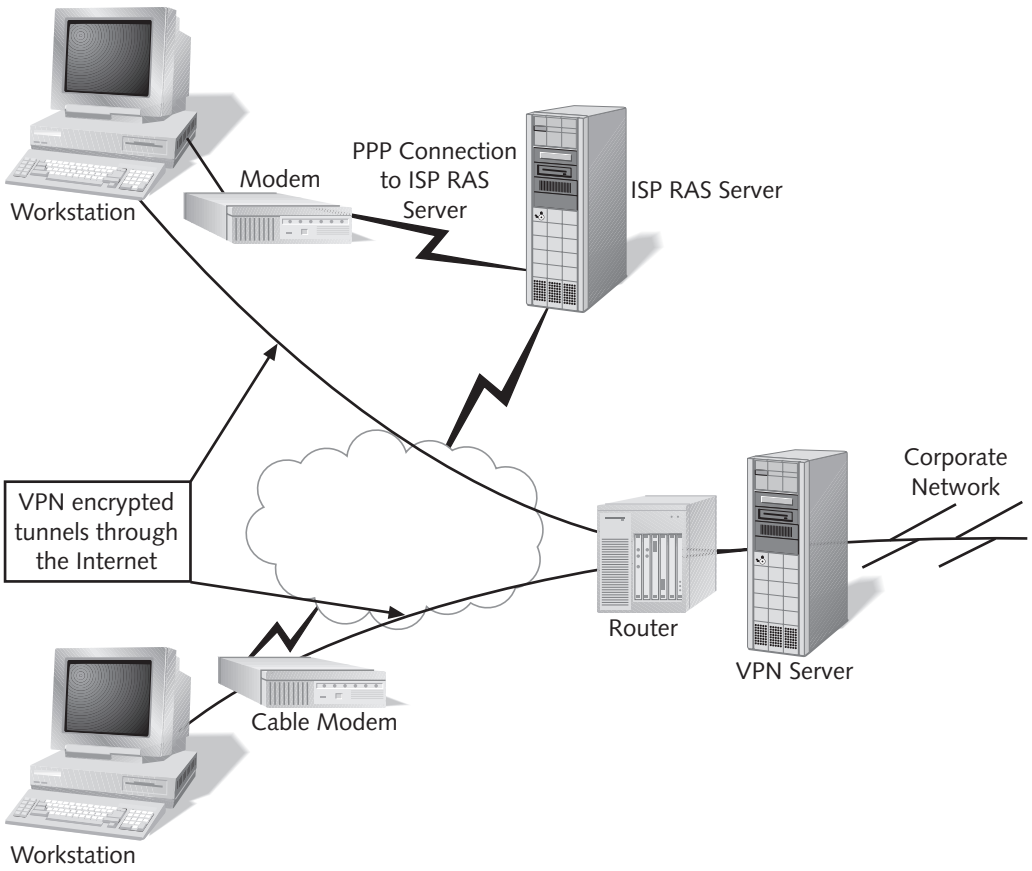

Figure 8-8 Creating a Virtual Private Network connection

Chapter 9, "Securing Access Between Corporate Locations" will go into much more detail on how virtual private networking actually works. It will include specifics on the VPN protocols supported by Windows 2000 and how to use the VPNs to create secure connections through the Internet between two corporate offices.

VPN Server Configuration

In many ways, configuring both the server and client for VPN is similar to configuring them for remote access. From the client, the VPN connection is configured as a network and dial-up connection. From the server, the **Routing and Remote Access Service (RRAS)** is used to configure the VPN server.

Windows 2000 supports two VPN protocols: Point-to-Point Tunneling Protocol (PPTP) and **Layer Two Tunneling Protocol (L2TP)**. Configuring the VPN server is similar for either protocol. The first step is to install two network adapters in the VPN server. One of the adapters must have a permanent connection to the Internet, and the other adapter is connected to the internal corporate network. Each of the adapters should be installed and tested before continuing.

After installing and testing the network adapters, follow these steps to configure the VPN server.

1. Configure the network adapter connected to the Internet with a valid IP address, subnet mask, and default gateway. Your ISP will provide you with the correct information.
2. Configure the network adapter connected to the corporate network with a valid internal IP address, subnet mask, and name server address (that is, DNS or WINS). The card should not be configured with a default gateway.
3. Choose **Start/Programs/Administrative Tools/Routing and Remote Access**.
4. Right-click the **RRAS server**, and click **Properties**.
5. Verify that the Router and the Remote access server check boxes are selected. See Figure 8-9.
6. Click the **IP** tab, and verify that Enable IP Routing is selected. Select whether the clients will receive IP addresses from a static pool or from DHCP on the network. Click **OK**.
7. Expand the RRAS server container, right-click **Ports**, and then click **Properties**.
8. Verify that WAN Miniport (PPTP) is selected and choose **Configure**. See Figure 8-10.
9. Select **Remote access connections (inbound only)** *and/or* **Demand-dial routing connections (inbound and outbound)**. The choice of which (or both) selection(s) will depend on your exact use of the VPN server. Specify the maximum number of PPTP ports available. Configure the server with enough PPTP ports to make sure that a port is available for each concurrent client that may connect.

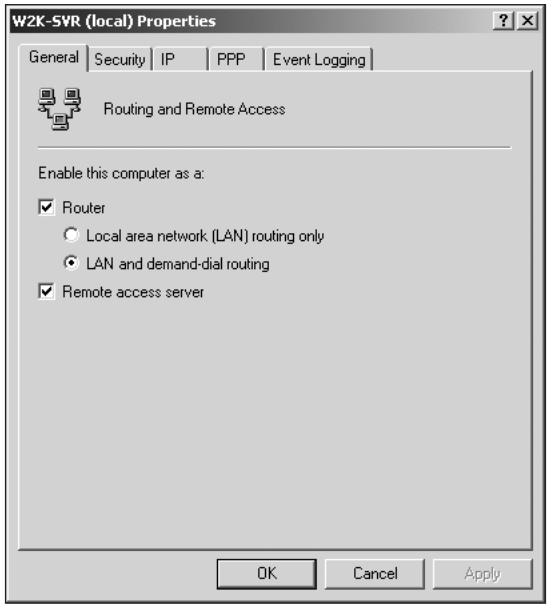

Figure 8-9 Configuring RRAS as a VPN server

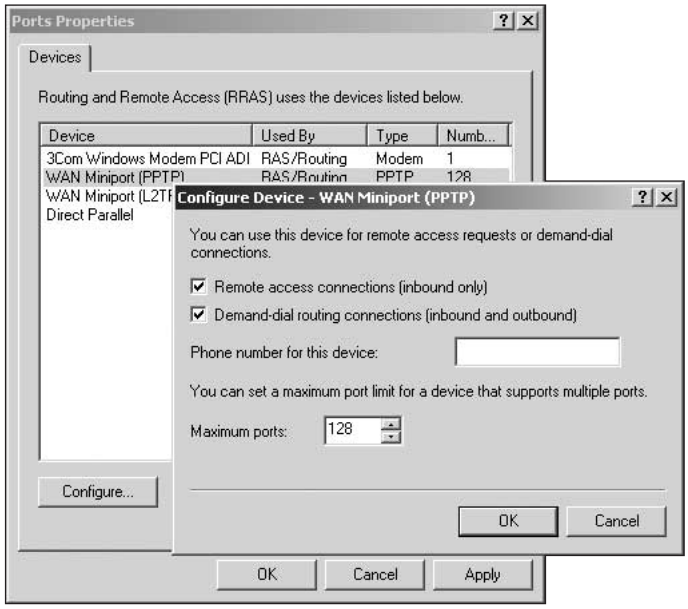

Figure 8-10 Configuring the VPN port settings

10. If you also have clients that are connecting using L2TP, repeat the port configuration for WAN Miniport (L2TP).
11. Click **OK** twice to accept the changes.

VPN Client Configuration

Configuring a VPN client is similar to configuring the dial-up client. To configure the VPN client, use the following procedure:

1. If you are connecting to an ISP using a dial-up account, create the dial-up connection first, using the procedure described above. If you are using a permanent connection to the Internet, such as a cable modem or DSL connection, you do not have to configure an initial connection.
2. Click **Start**, point to **Settings**, and click **Network and Dial-up Connections**.
3. Double-click **Make New Connection**.
4. The Network Connection Wizard starts. Click **Next**.
5. You are given the option of creating several different types of network connections. Choose **Connect to a private network through the Internet**. Click **Next**.
6. You are given a choice of whether you want to first dial an initial connection before creating the VPN. See Figure 8-11. If you are using an ISP dial-up account to access the Internet, select the dial-up connection name. If your connection to the Internet is always active, select **Do not dial the initial connection**. Click **Next**.
7. Type in the IP address or host name for the VPN server to which you are connecting. Click **Next**. (Note: If you do not have a dial-up adapter or modem installed, you may not get this dialog box.)
8. Choose whether to configure the connection for all users on the computer. Click **Next**.
9. Choose whether to share the connection with other computers on your network using Internet Connection Sharing. Click **Next**, and then click **Finish**.

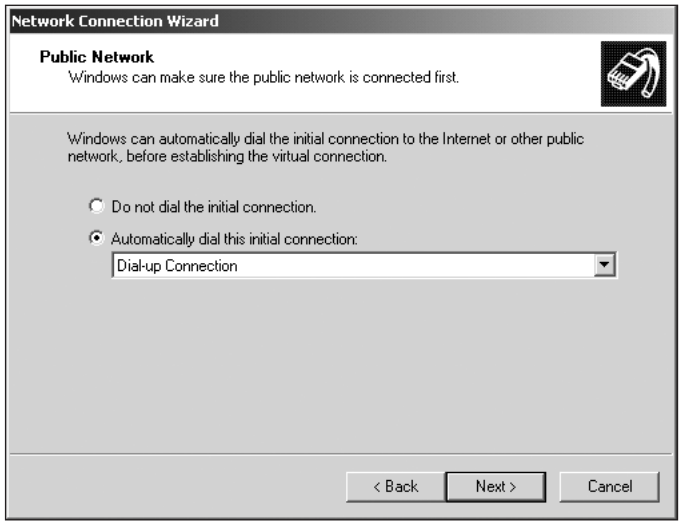

Figure 8-11 Configuring an initial connection for the VPN client

SECURING REMOTE ACCESS

One of the essential components of your security plan will be how to maximize the security of remote access to your network. As mentioned earlier, remote access is probably the only access point to your network other than your Internet connection; therefore, it must be made as secure as possible. The following sections discuss the options to enhance the security of remote access.

Configuring Remote Access Authentication

The first step to making remote access as secure as possible is to configure the remote access authentication. Windows 2000 Server provides several options that can be used to authenticate users who connect to a remote access server using either dial-up or VPN. The default protocols, **Microsoft Challenge Handshake Authentication Protocol (MS-CHAP)** and **MS-CHAP version 2** are sufficient for most situations if the client computers are all using Microsoft Windows 95 or later, NT 4.0, or Windows 2000.

MS-CHAP

MS-CHAP uses a challenge and response handshake for authentication. To begin the authentication process, the RRAS server sends the client a challenge that includes a challenge string and session identifier. The client responds with a username, session identifier, password, and a nonreversible encryption of the challenge string that it received from the server. The server will check and authenticate the response if it is valid.

MS-CHAP is quite secure because the password is never sent on the network in an unencrypted form.

MS-CHAP version 2

MS-CHAP version 2 is similar to MS-CHAP, but with enhanced security options, including the following:

- LAN Manager authentication is not supported, thus improving security.
- The cryptographic key used for encrypting data is based on the user's password and the challenge string, rather than on the password only.
- Two-way authentication is supported, which means that the client can also authenticate the server.
- A different key is used to encrypt data each time a user connects, and a different key is used to send and receive data.

During MS-CHAP version 2 authentication, the RAS server sends a challenge that includes a session identifier and challenge string. The client responds with the username and a peer challenge string that it generates. The client also sends a one-way encryption of its peer challenge string, the server challenge string, the session identifier, and the user's password. The server checks this response, then informs the client if authentication was successful and sends a response based on its challenge string, the peer challenge string, the client's encrypted response, and the user's password. The client must now authenticate the server's response. The client will not use the connection if unable to authenticate the server's response.

MS-CHAP (version 1 or 2) is needed for password changes during logon and for using Microsoft Point-to-Point encryption (MPPE) for PPP or PPTP encryption.

Extensible Authentication Protocol (EAP)

Windows 2000 supports **Extensible Authentication Protocol (EAP)** as a remote access authentication protocol. EAP is different from the other options in that it is really a standard for designing other authentication processes. Two options are currently supported for Windows 2000 authentication: EAP-Message Digest 5 Challenge Handshake Authentication Protocol (EAP-MD5 CHAP) and EAP-TLS. EAP-MD5 CHAP uses challenges and responses sent as EAP messages between the client and server. It is typically used to support authentication based on usernames and passwords. EAP-TLS is used with certificate-based security. For example, you would use EAP-TLS if you were using smart cards for logon and authentication. EAP-TLS is supported only if your RAS server is part of an Active Directory domain.

EAP also provides an Application Programming Interface (API) that supports the creation of any device or process to perform the authentication. This means that EAP can be used when developing a new authentication process, such as a retina scan, simply by designing the new process to use the EAP APIs.

Challenge Handshake Authentication Protocol (CHAP)

CHAP is an industry-standard challenge-response authentication protocol similar to MS-CHAP. You will need to use CHAP for the highest level of security if you are supporting dial-up clients that are not running Microsoft operating systems.

Shiva Password Authentication Protocol (SPAP)

You will need to use SPAP when connecting a Windows 2000 system as a client to a Shiva LAN Rover or when allowing Shiva clients access to your RAS server. SPAP uses the same encryption key every time a user connects to the server, so it is susceptible to replay attacks.

Password Authentication Protocol (PAP)

The password is sent on the network in clear text, so it is insecure. PAP should be used only in cases in which the client cannot support any other option.

Unauthenticated Access

Windows 2000 Server supports remote access connections that do not require any authentication. This option uses the Guest account to gain access to network resources, so the account needs to be enabled, and strict permissions need to be set for the account.

If you are using a Windows 2000 RRAS server as the remote access server and Windows 95 or later Microsoft clients, you will usually accept the default of supporting only MS-CHAP version 1 and version 2 authentication. In a more secure environment, you can enforce the use of MS-CHAP version 2 only. (To use MS-CHAP version 2 with Windows 9x clients, you must upgrade the dial-up client.) The most secure option for remote access authentication is EAP-TLS because the user must not only know a password or PIN number, but must also have physical possession of the smart card.

Another option for securing remote access is to use a random password-generating device for both the client and the server. In this case, the user is given a small device that generates a unique password every few minutes. When the user tries to access the remote access server, the user must type in the password that is currently displayed on the device. The server uses the same algorithm to generate passwords, so it will have the same current password for the user, and the user will be authenticated. The advantage of this type of device is that it protects you in case an attacker manages to install an application on your computer that captures your keystrokes as you type in your password. Because the password is only valid for a very short time period, the attacker will not be able to use the password to log in again. This type of system requires dedicated third-party hardware and software to implement.

Configuring Authentication Options on a Remote Access Server

To configure the authentication options on a remote access server, use the following procedure:

1. Open Routing and Remote Access from the Administrative Tools folder.
2. Right-click the server name, and click **Properties**.
3. Click the **Security** tab.
4. You are given a choice of using either Windows or RADIUS for authentication and user accounting.
5. Click the **Authentication Methods** button. See Figure 8-12.

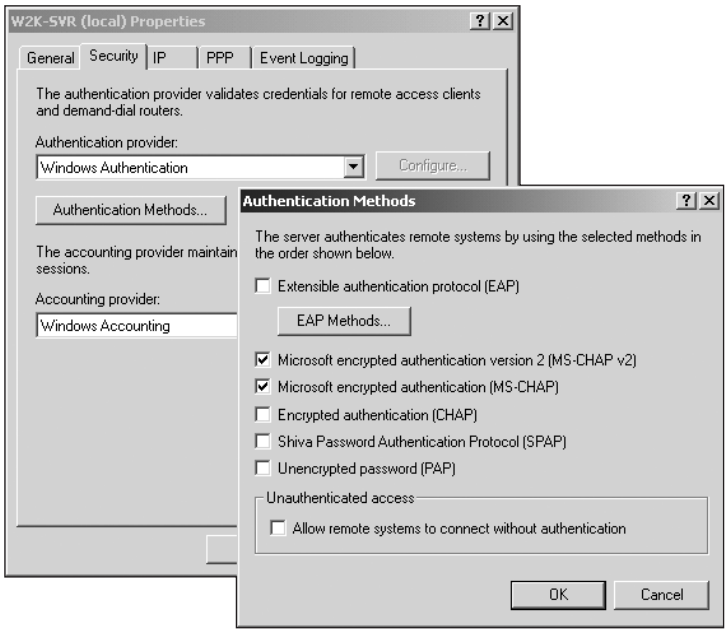

Figure 8-12 Configuring the RRAS server authentication method

6. Select the authentication options you want to use, and then click **OK**.

The authentication methods can also be configured in remote access policies.

Configuring Authentication Options for a Remote Access Client

To configure the type of authentication the remote access client will use, follow this procedure:

1. Open the remote connection properties in Network and Dial-up Connections.
2. Click the **Security** tab.
3. You can configure the option to use a secure or unsecure password or to use a smart card. To configure other options, select **Advanced (custom settings)** and click **Settings**. See Figure 8-13.

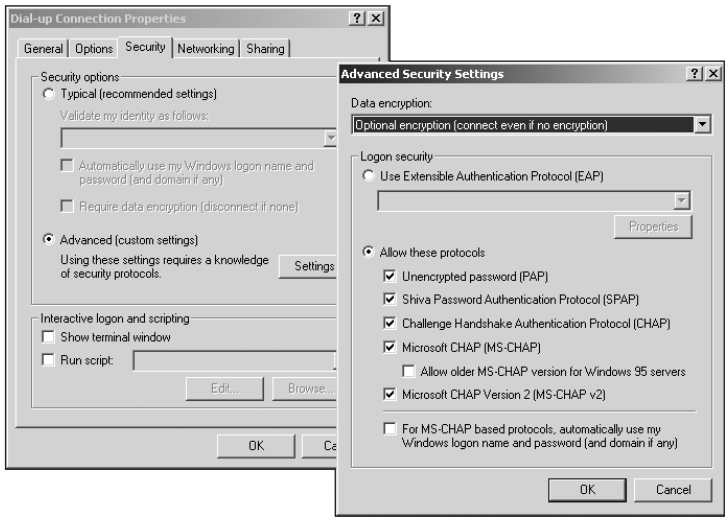

Figure 8-13 Configuring authentication options for a remote access client

4. Using the advanced settings, you can configure whether to used data encryption for all connections, as well as what types of authentication protocols to use. Click **OK**.

Configuring Callback Options

In addition to the authentication methods available in Windows 2000 Server, you can also use callback and caller ID to enhance security on your RAS server. **Callback** enhances network security because you can control the telephone numbers from which a dial-up connection can be made. If you implement callback, the client will call the RRAS server and go through the authentication process. Once the user is authenticated, the RRAS server will break the connection, and then call the user back at a preset number. This option is useful when the user calls in from the same location all of the time because you can limit the locations that the user can call from.

With **caller ID**, you can configure the dial-up connection so that a user can dial in from only one specific phone number. For caller ID to work, the entire phone system, from the user to the Windows 2000 remote access server, must support caller ID. If any part of the system does not support caller ID, then the connection will fail. Caller ID can also be used for a VPN connection, in which case the caller ID is set to a specific IP address.

To configure callback and caller ID dial-up security in Windows 2000, use the following procedure:

1. Open Active Directory Users and Computers and locate the user account.
2. Right-click the account name and click **Properties**.
3. Click the **Dial-in** tab. See Figure 8-14.

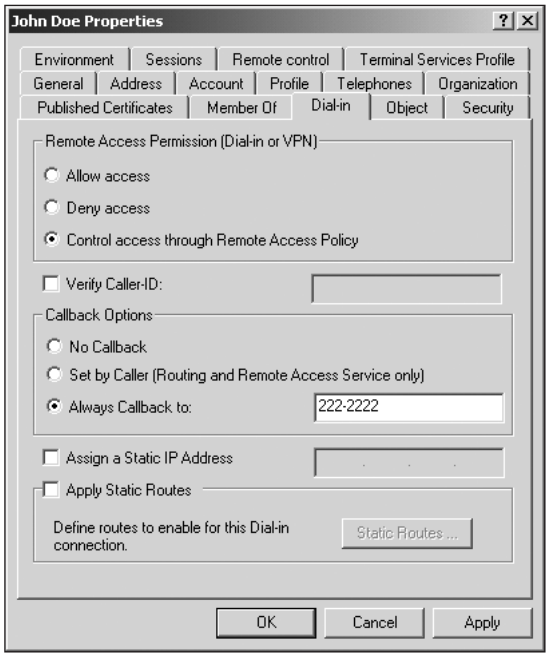

Figure 8-14 Configuring the callback options on a user account

4. Select the **Verify Caller-ID** check box, and then type the number or IP address from which the user will connect.

The Verify Caller-ID check box is only available if the domain is configured for native mode operation or if the account is a local account on a standalone computer.

5. Select *one* of the following callback options:
   a. No Callback—The server will not call the user back.
   b. Set by Caller—After the user has been authenticated, the user can supply the number for the server to call back.
   c. Always Callback to:—The server always calls back to the preset number.

Remote Access Account Lockout

Another option that can be used to increase remote access security is remote access account lockout. The account lockout feature is used to lock out user accounts if too many attempts have been made to connect to the server using an incorrect password. This is especially important for VPN connections, because the connection will always be available on the Internet and accessible to attacks. By setting the limit for how many incorrect passwords are allowed, you can prevent password-guessing attacks on your remote connection.

Account lockout is normally configured through the account policies on the domain group policy. However, account lockout for remote access can be administered separately to create different settings, depending on whether the logon attempt is on a LAN connection or on a remote access connection. The two options are compatible, however. If the maximum number of incorrect passwords is set in both the domain policy and for remote access, then the setting with the lowest number of retry attempts will be applied. For example, if the domain policy sets the maximum number of password retries at three, and the remote access policy is set at 5, then the account will be locked out based on the domain policy before it is locked out using the remote access setting.

By default, remote user connection attempts are locked out with the same setting as the domain setting, so if you want the password lockout policy to be the same for both LAN connections and remote connections, then configure the setting just in the domain policies. However, if you want to set the remote access lockout policy at a lower number than the LAN policy, you must configure the remote access setting on the RRAS server. To enable user account lockout for remote access, edit the following Registry key:

```
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\
RemoteAccess\Parameters\AccountLockout\MaxDenials
```

By default, this key is set to 0 (zero), which means that account lockout specifically for remote access is not enabled. To lock an account out after five attempts, set the value to 5.

To configure how long Windows 2000 Server should wait to reset the failed attempts counter, edit the following Registry key:

```
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\
RemoteAccess\Parameters\AccountLockout\ResetTime
```

The default setting for this key is 0xb40 or 2880 minutes (48 hours). This means that if the MaxDenial setting is reached within 48 hours, the account will be locked out. A successful logon resets the failed attempts counter.

To reset an account that has been locked out for remote access, delete the following key:

HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\
RemoteAccess\Parameters\AccountLockout\DomainName:username

Administering a separate account lockout policy for remote access users takes a significant amount of administrative effort, particularly since you must edit the Registry to enable the separate policy, and you have to delete a Registry key to unlock an account that has been locked out using the remote access setting. Because of this, you may want to configure the default setting on the domain Group Policy to take into account remote access as well. If you would like to lock out user accounts after five attempts on a LAN connection, but would like to have remote access users locked out after three attempts, it is much easier to compromise at four attempts for both connections rather than manage the Registry changes. However, if you work in an environment that requires high security, and you have a VPN connection that is accessible from the Internet, you may need to set a lower limit for bad logon attempts for the remote access users. Be sure that you have the administrative resources to manage resetting the account lockouts.

User Education

Some of the most insecure components of your remote access solution are the people who use remote access. In many cases, these users do not understand the technology that they are using and may not be aware of the security risks involved in remote access. Some of the users are probably quite frustrated with the low speed and unreliable nature of remote access, especially if they are using a dial-up connection. This combination of lack of knowledge and a high level of frustration makes remote access users susceptible to social engineering attacks.

As discussed in Chapter 1, "Identifying Security Risks," a **social engineering attack** is an attack aimed at the users of a technology rather than at the technology itself. The most common way to launch a social engineering attack is to contact the user using false credentials. For example, the remote access user may get a phone call in which the caller identifies himself as working on a project for the IT department to enhance the reliability of the dial-up remote access, and, as part of the testing, the caller needs to connect to the dial-up server using the client's username and password. Or the client may be asked to reset her password to a specified value, again for testing purposes. The attacker may ask the user to explain the procedure she uses for connecting remotely and, as part of that explanation, may ask for the username and password that the client is using. The social engineering attack can be very ingenuous. The attacker may refer to the head of the IT department by name, as the sponsor for the project, or may mention other people on the help desk as also working on the project. If the attacker has some information about the user, such as the user's department or supervisor, he may use this to enhance his appearance of legitimacy. The entire goal of the social engineering attack is to get the user to trust the attacker to the point where the user does what the attacker asks.

The only defense against a social engineering attack is user education. As part of your security plan, you should very clearly identify acceptable practices for remote access users. One of these acceptable practices is that they should never, under any circumstances, reveal their remote access password to anyone on the phone or in response to an e-mail. If your company really is testing remote access connectivity, then the people doing the testing should use test user accounts, or if they need to use a specific user account, they should reset the password for the account to a known value, and then tell the user what their new password is. Your security policy should also define a specified procedure for the remote access users that they can use to confirm any type of request for user information. (Actually the policy should apply to all users, not just remote access users.) If a person receives a call from someone who says they are working at the help desk, the policy should outline what types of information the user can provide over the phone, and also provide the user with some way of validating the identity of the caller. For example, you may want to implement a policy that requires all users to disconnect from the phone call and call the help desk back at a specified number.

Another user education and policy component with remote access is the administration of user accounts that have remote access. In most companies, the default policy when creating a user account should be to deny them remote access. The user should be given remote access only if the user has a clear need for it. Even more importantly, if a user leaves the company, their user account should be immediately disabled so that they cannot get access to company data remotely or locally. In addition, you should audit all remote access to your network, and monitor the log files closely.

Every company has different security requirements for remote access, but one of the key points to keep in mind is that the remote access connection is one of the most important security risks. If an attacker can get access to a user's logon name and password, and you don't have policies such as callback in place, the attacker can log onto your network from anywhere. If the attacker logs in after regular work hours, he may have many hours of access to your network before anyone discovers what is happening. In a worst-case scenario, if the attacker gets access to an account with administrative permissions, he may steal or destroy a great deal of confidential information. The attacker may be able to launch an application on your network that captures other user's passwords and forwards them to an Internet address.

Remote Access Policies

One of the biggest improvements in Windows 2000 for supporting remote access is the introduction of **remote access policies** to manage the configuration of client access. In previous implementations of RAS, the only option for configuring user access was to manage each user individually, which meant that you had to access each user account to

allow or deny dial-up access. If you had a large number of users that were dialing in, administering the user accounts individually was a considerable amount of work. As well, the options for configuring remote access were very limited. The options for configuring remote access on a Windows NT 4.0 RAS server were also very limited. For instance, you could only allow or deny access and configure callback options. Remote access policies in Windows 2000 provide much better functionality. For example, you can now set remote access policies for Active Directory security groups rather than on an individual basis. In addition, you can configure options such as the hours that a person can connect remotely to the network and what protocols the user can use to connect.

Remote Access Policy Concepts

By default, the remote access settings for individual users can still be configured for each account in Active Directory Users and Computers. When you access the user dial-in properties, you are given the choice to Allow access, Deny access or Control access through remote access policy.

The option to set remote access based on remote access policies is available only on the user account in Active Directory Users and Computers when the domain is configured to be running in native mode. However, remote access policies are always evaluated as part of the dial-up process.

If you configure a remote access policy, whether a user can dial-in depends on conditions, permissions, and profile settings. A condition is a setting, such as which phone number the person dials from, what Active Directory group the user belongs to, or what time of day the user can dial-in. When a user tries to connect using remote access, RRAS tries to match the conditions set in a remote access policy to the remote connection. For example, if the user is a member of the Managers group, and the person is trying to connect at 11:00 PM, RRAS tries to locate a remote access policy that matches these conditions.

The permissions on the remote access policy are based on the Active Directory dial-in settings for the user account and on the permissions assigned by the remote access policy. If the remote access conditions are met, RRAS then checks Active Directory to determine the user permissions. If the user is denied dial-in permissions, then the user will not be able to connect. If the user is allowed dial-in permissions, then the policy evaluation will continue. If the user permissions are controlled through the remote access policy, then the permissions on the remote access policy are checked. The remote access policy permissions either allow or deny access.

If the user has been granted remote access based on the evaluation of conditions and permissions, the last step is to evaluate the profile settings. The profile setting includes options such as disconnection settings, IP address assignments, authentication settings, and multilink settings. Again, the connection attempt is evaluated to determine whether the user will be able to connect. For example, the profile might require that the user

must authenticate using MS-CHAP version 2. If the client cannot do this, the connection will be denied.

To understand how the conditions, permissions, and profile settings work together, consider a situation in which a user named Jane Doe, a member of the Sales group, tries to connect to the RRAS server using a VPN that is configured with the remote access policy shown in Table 8-1:

Table 8-1 Remote Access Policy conditions, permissions, and profile

Conditions	Dial-up is supported from Monday to Friday, 6:00 AM to 12:00 AM only. The user must be a member of the Sales or the Managers groups. The user must use either PPTP or L2TP tunneling protocols.
Permission	The dial-up permissions for Sales1 are configured to Control access through remote access policy, and the callback option is set at No Callback. The permissions on the remote access policy allow access.
Dial-in Profile	Users must use MS-CHAP version 2 to authenticate. Users must use strong encryption for this connection. The user can stay connected for no more than 15 minutes. The user can only connect using a VPN.

When the user connects to the server, the conditions for this policy are examined to see if the user meets them. All of the conditions must be met in order for Jane Doe to connect to the server. If she does not meet all conditions, the server will look for another policy and check the conditions for that policy. The policies are processed in the order that they are listed in RRAS, and the first policy in which the user meets all of the conditions is the policy that is applied. In this example, if the time is correct, and Jane Doe is connecting using PPTP or L2TP, the conditions are met. The user's permissions are then checked. The permissions are a combination of the user account permissions and the remote access policy permissions. For example, if the Sales group were given permission to dial-in, but Jane Doe was denied permission, then she would not be able to connect. In this case, because Jane Doe's permissions are controlled by the remote access policy and the policy settings allow access, she will be able to connect. Then, the dial-up profile is checked and, again, all of the settings must be met before the user is provided access. If Jane Doe's computer supports MS-CHAP version 2 and strong encryption, she will be able to connect and stay connected for 15 minutes.

Configuring Remote Access Policies

Using all of these options gives you a great deal of flexibility in configuring user dial-up access. To create and configure remote access policies, use the following procedure:

1. Open **Routing and Remote Access** from the Administrative Tools folder.

2. Expand the Server container, right-click **Remote Access Policies**, and choose **New Remote Access Policy**.
3. Type a name for the policy and click **Next**.
4. Click **Add** to add the conditions that will be used for this policy. See Figure 8-15. The conditions that you configure are evaluated every time a user tries to connect to the RAS server, and the user is given remote access only if the conditions on one of the remote access policies exactly match the user's situation.

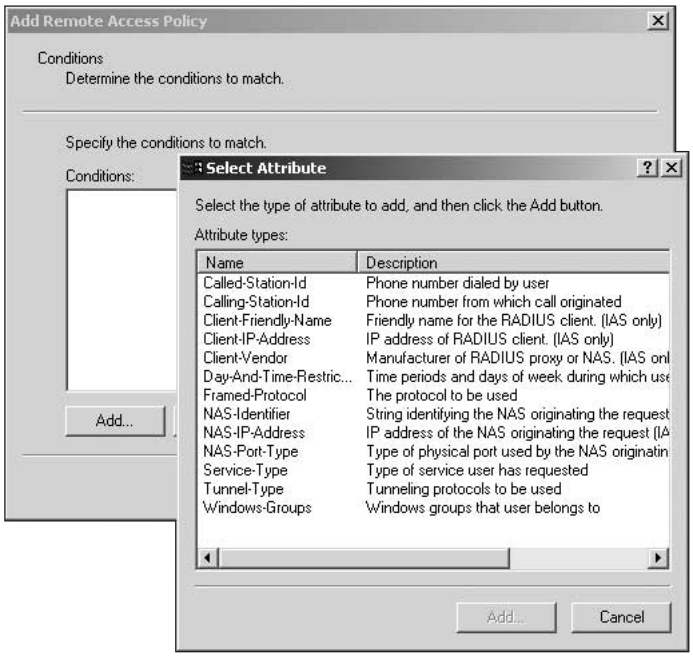

Figure 8-15 Configuring the remote access policy conditions

5. One of the conditions you might frequently use is the Windows-Groups condition, which can be used to designate which Active Directory group will use this connection. To add this condition, select **Windows-Groups** and click **Add**. Click **Add** again and select the groups that should be included in this policy. When you have added all the groups, click **OK**.
6. Another option you might want to configure is the time of day that the group is allowed to dial-in. To do this, click **Day-and-Time Restrictions** and click **Add**. Select the time of day that the users should be granted access, and then click **Permitted**. See Figure 8-16. Click **OK**. When you are finished selecting the conditions, click **Next**.

7. You are then given a choice whether you want to grant remote access permission or deny remote access permission to those users who match the conditions. This is a powerful option in remote access policies. For example, you can choose to deny permission using this policy. Then, when a user tries to connect, and the user conditions meet the conditions specified in the policy, the user will be denied access. If a policy denies a user access, RRAS will not check any more policies, but will deny access. Once you have configured the permissions, click **Next**.
8. You are then given the option to edit the profile. To edit the profile, click **Edit Profile**. Otherwise, click **Finish**.

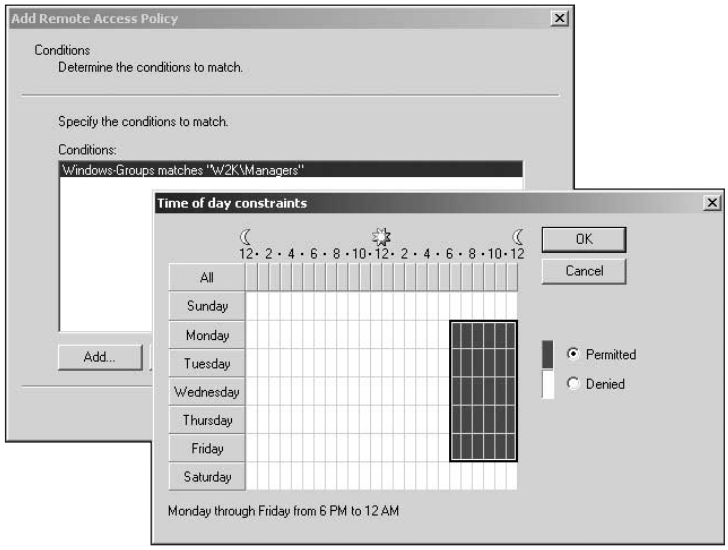

Figure 8-16 Configuring time-of-day constraints for a remote access policy

If you choose to edit the profile, you are presented with the Edit Dial-in Profile dialog box. See Figure 8-17. You can configure many different properties on each profile to control the settings for this remote access policy. The options on each tab are discussed in Table 8-2.

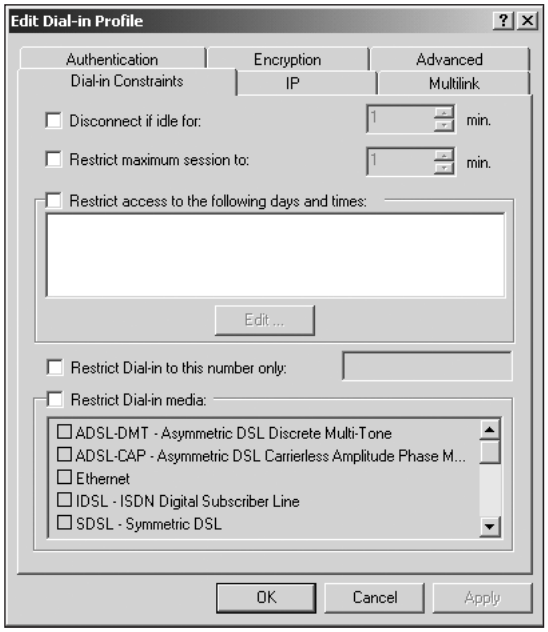

Figure 8-17 Configuring the dial-in profile

Table 8-2 Configuring the Dial-in profile

Dial-in Constraints	Used to configure disconnection times, maximum length for a connection, time restrictions, dial-in phone number restrictions, and what type of media can be used.
IP	Used to configure how the client will get an IP address and IP address filtering, which can then be used to control what types of protocols the client can use.
Multilink	Used to configure whether or not the users for this profile are able to use multilink and, if they use multilink, how many connections they can use and at what point one of the connections will be dropped if there is not enough traffic on the connections. Bandwidth Allocation Protocol (BAP) is also enabled on this page.
Authentication	Used to assign authentication methods that will be accepted for users of this profile.
Encryption	Used to assign the level of encryption to users of this connection. The options are: No Encryption, Basic (IPSec 40-bit DES or MPPE 40-bit), or Strong (IPSec 56-bit DES or MPPE 56-bit).
Advanced	Used to set other RADIUS-specific or hardware vendor-specific settings, which can then be used to limit access for users of this profile.

To edit the conditions after the policy has been created, right-click the policy in RRAS and click Properties.

After a remote access policy has been set up, the administrator must configure each user account in Active Directory Users and Computers to use a remote access policy to control user permissions.

Planning for Remote Access Policies

Remote access policies are a powerful new tool for you to use to control remote access to your network, and as with all powerful tools, you need to plan your implementation carefully. Before you implement remote access policies, you need to test your configuration thoroughly. In particular, you need to include the following components in your remote access policy planning:

- *Determining the remote access policy order in RRAS*—The order in which the policies are listed in RRAS determines how they are applied. When a user tries to connect to the RRAS server, the server begins with the first policy in the list and checks the conditions. If the conditions match the connection attempt, then the server checks the permissions and the profile. If the conditions do not match the connection attempt, then the server checks the rest of the remote access policies in order until a policy that matches the conditions is found, or all the policies have been checked. You need to plan the order of the policies carefully. For example, you may create one policy that allows the members of the Managers group to remotely access the network from 6:00 PM to 12:00 PM and another policy that explicitly denies the Domain Users group permission to connect remotely at any time. If the managers' policy is evaluated before the Domain User policy, the managers will be able to gain access. If the Domain Users policy is listed first, the managers, who are also part of the Domain Users group, will not be able to access the server.
- *Using the default policy*—When RRAS is installed, a default remote access policy (called Allow access if dial-in permissions is enabled) is created. This policy denies everyone permission to connect remotely, unless they have been given explicit permission to dial in in Active Directory. In most cases, you should leave the default policy in place as the final policy in RRAS. This means that you can define all your policies that grant permission for users to connect, and if none of the policies apply, they will be denied access.
- *Using policies to deny access*—You can use remote access policies to deny user permission to connect remotely. In most cases, you would use this option to set exceptions to a larger rule. For example, you may want to grant the All Managers group permission to dial-in, but you do not want to give the Admin Managers this permission, even though the Admin Managers are also in the All Managers group. In this case, you could create a policy to deny the

Admin Managers dial-in permission, and another policy to grant the All Managers group permission. As long as the Admin Managers policy is evaluated first, they will not be able to dial-in, while the other members of the All Managers group will.

Document all policies thoroughly, so that if a user is having trouble connecting to the RAS server, you can quickly identify which policies apply to that user.

Another issue that you need to plan for as you implement remote access policies is that remote access policies are stored locally on each RAS server. This means that if you are using more than one server for remote access, you must re-create each policy on each individual server. One of the ways to get around this limitation is to use **Internet Authentication Server (IAS)**, which is Microsoft's implementation of the Remote Authentication Dial-in User Service (RADIUS). With IAS, you can store the remote access policies on the IAS server and have all of the RRAS servers check the IAS server for the policies.

INTERNET AUTHENTICATION SERVER

If you are supporting a large number of remote access clients connecting to multiple RAS servers, administering all of these servers will be difficult even if you use remote access policies. One of the factors that makes this complicated is that the remote access clients might be connecting from many different places, including home offices, hotel rooms anywhere in the world, or through a VPN from a customer's network. In addition, the remote access servers may be located in many different locations, perhaps throughout the world. To support this type of environment, Windows 2000 includes IAS.

Introduction to Remote Authentication Dial-in User Service (RADIUS)

RADIUS is an open standard that is widely used for remote access authentication in large, distributed environments, such as ISPs. The purpose of RADIUS is to centralize the management of remote access users and policies in an environment in which the remote access servers might be widely distributed. In a RADIUS implementation, one server contains all of the user accounts and remote access permissions, and multiple dial-up servers are configured to forward authentication requests to that one RADIUS server. This centralizes the management of the user accounts and policies, while still distributing the load of providing the dial-up access for users.

RADIUS can be implemented in many different configurations. One of the simplest implementations is to have a single server that supports the modems that are used for the remote dial-up, and another single RADIUS server.

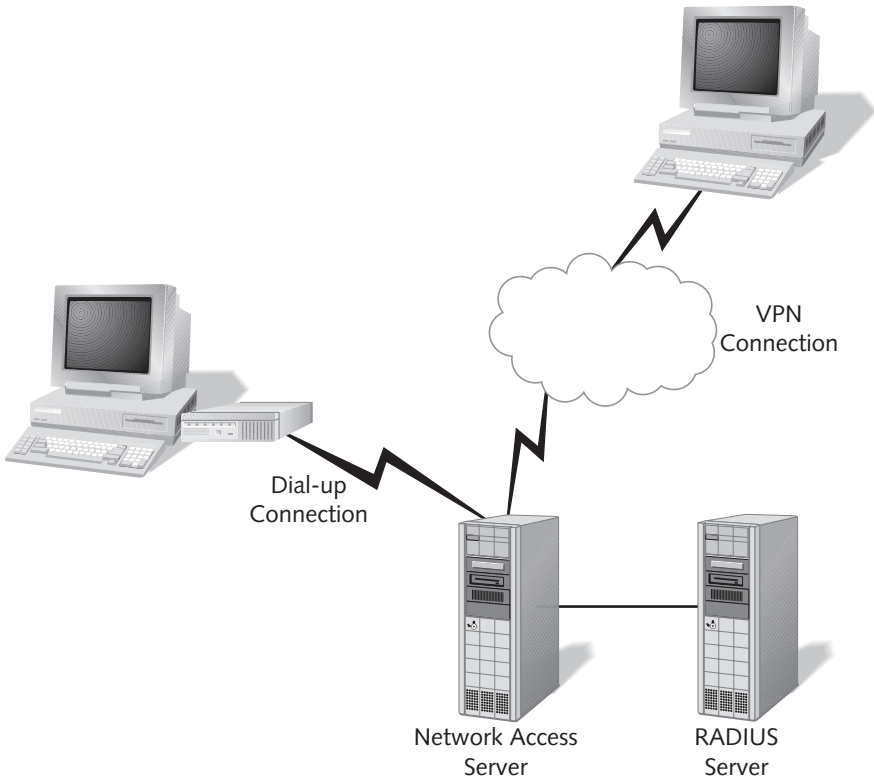

Figure 8-18 A simple RADIUS server implementation

RADIUS uses some specific terminology to describe the different server roles in this environment. The server that supports the modems is called a **Network Access Server (NAS)**. The NAS is also a RADIUS client because it acts as the client in the authentication process by forwarding all authentication requests to the RADIUS Server. The RADIUS server has a database of users and policies that are checked when the server receives an authentication request. Figure 8-18 illustrates a simple RADIUS server implementation.

RADIUS can be scaled to almost any scale and almost any remote access scenario. There are multiple RADIUS servers that have the same users and policy information in their database. Multiple RADIUS servers can be configured in this way for redundancy and load balancing. All of the other remote access servers are RADIUS clients. When they receive a remote access request, they forward the authentication to the RADIUS servers. Figure 8-19 shows what a complex RADIUS implementation might look like.

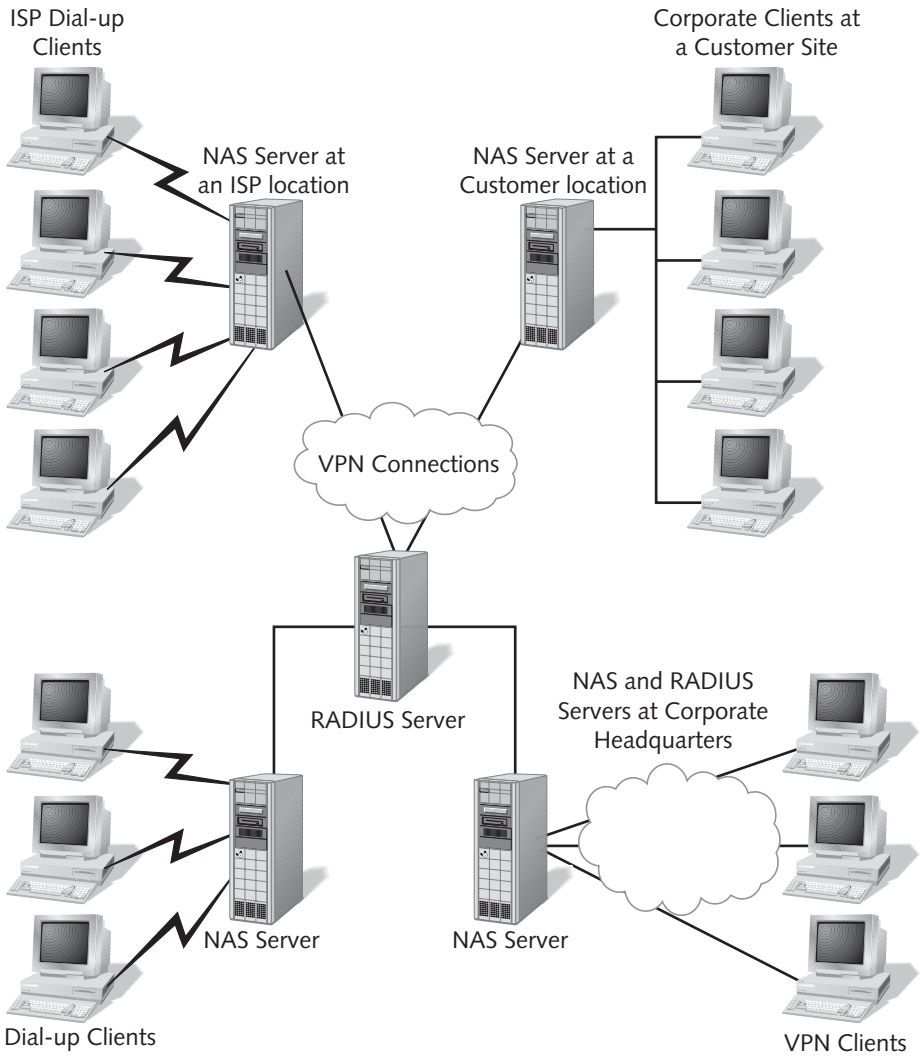

Figure 8-19 A complex RADIUS server implementation

This example shows some of the power of distributing the network access servers, while centralizing the remote access permission management. You can configure multiple NAS on your corporate network to forward requests to the RADIUS server. You can also configure a RADIUS client through an ISP. For example, if your company has a large number of users in a particular city, you may want to negotiate with an ISP in that city to host a RADIUS client for you. The ISP could configure one of their dial-up servers to forward all of the authentication requests to your RADIUS servers. The authentication would occur when the user dials the ISP, but you would control the user database on the corporate RADIUS servers. A similar configuration is possible if you

have a large number of users that are working for an extended period of time at a customer site, and you want to use a VPN from the customer site to the corporate office. You can configure a RADIUS client to authenticate all of the VPN connections at the customer site.

Windows 2000 Server can be integrated anywhere into the RADIUS configuration. RADIUS is based on an open standard, so all RADIUS implementations that adhere to the standard should be compatible. Windows 2000 includes IAS, which is a fully compliant RADIUS server. Windows 2000 RRAS can be configured as RADIUS client, connecting to either a third-party RADIUS server or to IAS. Windows 2000 IAS can operate as the RADIUS server, and other third-party RADIUS clients can refer the entire client authentication to the Windows 2000 server.

Windows 2000 IAS may be RADIUS-compliant, but it is also integrated with other Windows 2000 components. For example, IAS can be configured to use Active Directory as its user database, so that IAS will refer all authentication requests to Active Directory. This means that you can still use Active Directory to configure user remote access permissions.

Another area of integration of IAS with Windows 2000 is the option to centralize the management of remote access policies. If you have multiple RRAS servers and do not use IAS, you will have to copy all of the remote access policies to each server. However, after implementing IAS, you can configure all of the RRAS servers to check with IAS for remote access policies, so you only need to maintain a single copy of the remote access policies.

Implementing IAS as a RADIUS Server

To use IAS, you must first install the IAS service, and then configure the IAS options. Finally, you must configure all of your RRAS servers to use IAS for authentication and as the source for remote access policies.

Installing IAS

IAS is not installed by default in Windows 2000. To install IAS, use the following procedure:

1. Open the Control Panel, click **Add/Remove Programs**, and then click **Add/Remove Windows Components**.
2. Click **Networking Services** and click **Details**.
3. Select **Internet Authentication Service** and click **OK**.
4. Click **Next**, and then click **Finish**.

Configuring IAS

Once you have installed IAS, you have to configure the service so that it meets your security requirements and accepts connections from other RADIUS clients. To

configure IAS, open Internet Authentication Service from the Administrative Tools folder. As shown in Figure 8-20, the IAS administration tool contains three containers:

- *Clients container*—Holds entries for all IAS clients.
- *Remote Access Logging container*—Holds the default log file location and configuration. You can access the log file's properties to change the level of logging that IAS will perform, as well as change the log file location and how often a new file is created.
- *Remote Access Policies container*—Holds the remote access policies that have been defined in IAS. Managing remote access policies from IAS is identical to managing them with RRAS, except that the policy in IAS will be applied by all IAS clients.

To configure the default properties for the IAS server, right-click Internet Authentication Service (Local) and click Properties. See Figure 8-20. On the Service tab, you can type a name for the server, as well as configure logging of authentication requests. On the RADIUS tab, you can configure the port numbers used for authentication requests and accounting requests. The Realms tab is used to configure realm conversions if you already have dial-up clients that have been configured to use a different realm than you are currently using. For example, if you have clients that are using alias@Widgets.com to authenticate, but your domain name is MegaWidgets.com, you can use this tab to replace the realm name for all clients using Widgets.com.

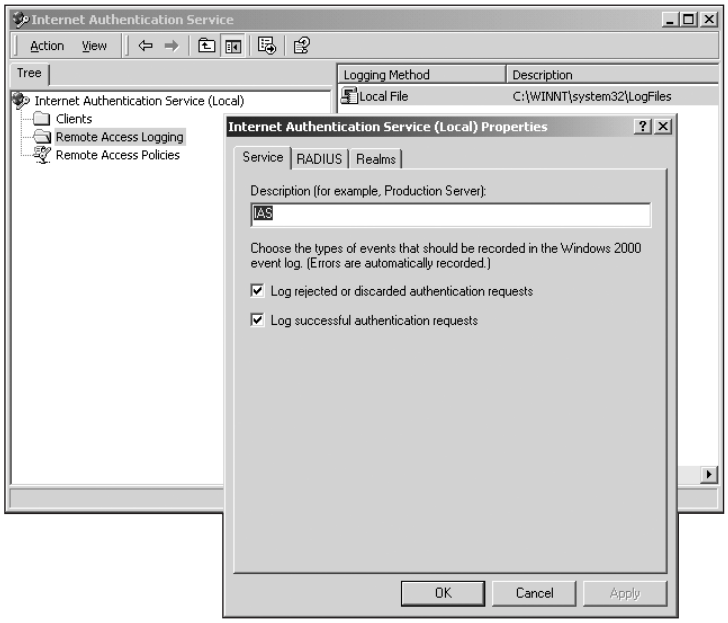

Figure 8-20 Configuring the Internet Authentication Service

If you want to configure the IAS server to check Active Directory for user information, you must register the IAS server in Active Directory. To do this, right-click Internet Authentication Service and select Register service in Active Directory.

Before you can configure your RRAS servers to use IAS for authentication and remote access policies, you must add the RRAS servers as IAS clients. To do this, use the following procedure:

1. Right-click the **Clients** container in Internet Authentication Service, and then click **New Client**.
2. Type a name for the RRAS server that will be acting as the IAS client. This name does not have to be the same as the actual computer name, but you may want to avoid confusion by using the same name. Click **Next**.
3. Type the IP address or DNS name for the IAS client and configure the standard that will be used by the clients. You can configure the Client-Vendor as Microsoft or accept the default of RADIUS Standard, as shown in Figure 8-21. Type the shared secret password. The shared secret is an important configuration option in IAS. The server is configured with this password, as well as all of the IAS clients, and this password is the primary security mechanism in IAS. You should make this password as long and complex as practical.

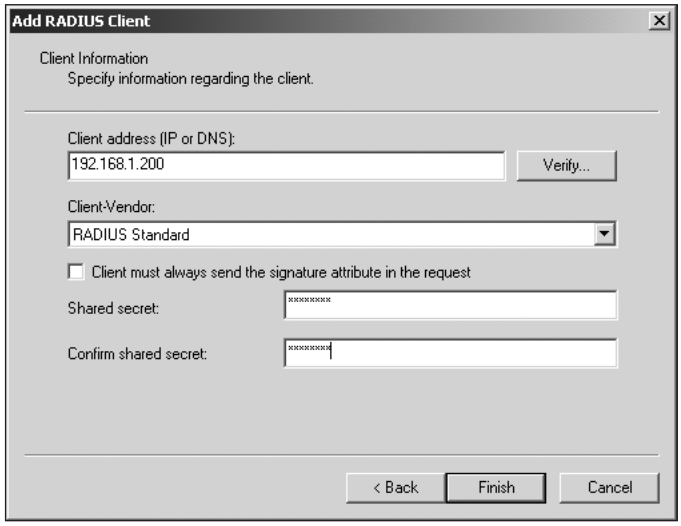

Figure 8-21 Adding a RADIUS client to IAS

4. After you have configured the settings, click **Finish.**

Configuring RRAS to Use IAS

The final step in configuring IAS is to configure the RRAS servers to use IAS to authenticate users, rather than Active Directory or the local user accounts. To do this, use the following procedure:

1. Open **Routing and Remote Access** from the Administrative Tools folder.
2. Right-click the server name and click **Properties**. Click the **Security** tab.
3. Configure both the Authentication provider and Accounting provider to use RADIUS.
4. Click **Configure** to configure the IAS server that RRAS will use.
5. Click **Add** to add the IAS server information. See Figure 8-22.
6. Type the IAS server name and configure the secret to be the same as the IAS server, and then accept the defaults for the other settings. Click **OK**. You also need to configure the IAS server for the other option on the Security tab.
7. You will receive a warning telling you to restart the RRAS service. Click **OK**, and then stop and restart the service.

You will have to repeat the last two procedures for all of the IAS client computers. Once this is complete, you can use IAS to configure all of your remote access policies.

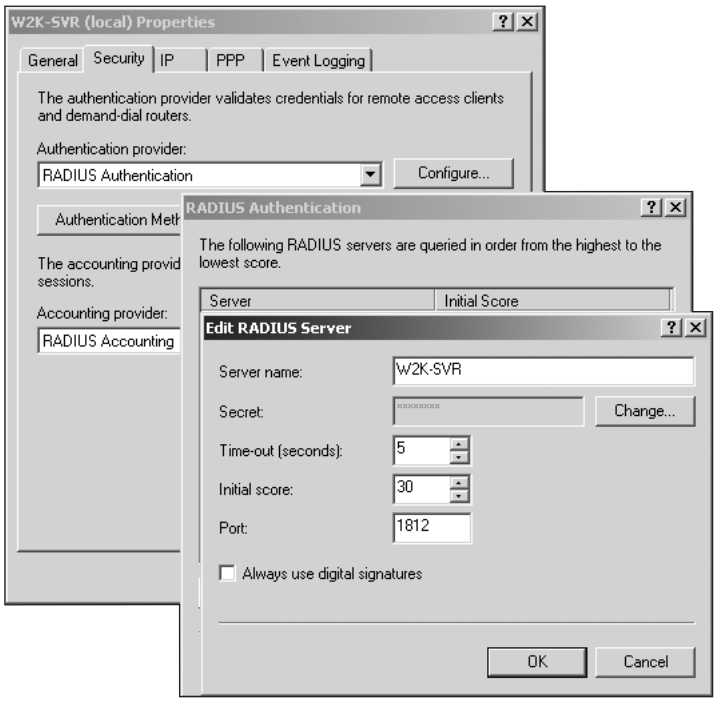

Figure 8-22 Configuring RRAS to use IAS

Planning Best Practices

- Remote access is one of the greatest security risks on your network. If your organization must provide remote access, plan the implementation very thoroughly.
- Windows 2000 can operate as a remote access server and support multiple modems in this configuration. However, if you require dozens of modems, you will need to buy a dedicated hardware device.
- If possible, force all clients to use MS-CHAP version 2 as their authentication protocol. This means that you will need to upgrade the dial-up client on the Windows 9x clients to the latest version. If you have nonWindows clients connecting remotely to your network, use CHAP authentication.
- Administering a separate account lockout policy for remote access users is difficult because you must edit the Registry to enable the separate policy, and you have to delete a Registry key to unlock an account that has been locked out using the remote access setting. To avoid this, use the default domain lockout policy for remote access as well.
- Any changes that you make to your remote access configuration must be tested thoroughly in a test environment before they are implemented in the production environment. Much of the remote access activity is probably taking place after regular business hours when the help desk might not be available. Therefore, your implementation of remote access should be as stable and tested as possible.
- Use remote access policies if you have more than a few users connecting to your network remotely. These policies provide a powerful tool to manage remote access, as well as enhance the remote access security. Implementing and testing remote access policies can be complicated in a large environment, but you will find that implementing them is well worth the effort.
- Very few companies are likely to need the functionality provided by IAS server because usually the remote access infrastructure is not that complex.

Chapter Summary

- Windows 2000 can operate both as a remote access client and as a remote access server. The remote access connection in either case can be through a dial-up connection using analog modems or ISDN lines, or through a Virtual Private Network connection. The Windows 2000 service that provides remote access is Routing and Remote Access.
- There are several ways to enhance the security of remote access in Windows 2000. These options include using the highest security authentication protocols (MS-CHAP version 2, or EAP), configuring callback and caller ID, and using

remote access account lockout policies. In addition to these technical solutions, you should also provide training and policies for remote access use so that the users can avoid social engineering attacks.

- One of the new features in Windows 2000 RRAS is the option to set up remote access policies to control who can connect to your network remotely. With remote access policies, you can configure remote access based on Active Directory groups, time of day, types of protocols used by the clients, etc.
- When a client tries to connect to a remote access server, and you are using remote access policies, a combination of conditions, permissions, and profiles is used to determine whether the user should get access. A condition has a setting, such as membership in an Active Directory group, or time of day. RRAS tries to match the conditions set in a remote access policy to the remote connection condition. If the user connection matches a condition, then the user permissions are checked. And finally, the user profile is checked. If the conditions are met, and the user is granted permission, and if the settings in the profile are met, then the user is granted remote access.
- Internet Authentication Service (IAS) is Microsoft's implementation of the RADIUS open standard. IAS enables you to have multiple RRAS servers in many different locations, but still manage your remote access database and settings from a single location. You can configure all of the RRAS servers to operate as IAS clients so that they will contact the IAS server when authenticating users and to access the remote access policies stored on the IAS server.

Key Terms

Bandwidth Allocation Protocol (BAP) — A protocol used to dynamically manage multiple links for remote access.

callback — The option available on a remote access connection in which you can configure the remote access server to call a user at a specified phone number after the client has connected and has been authenticated.

caller ID — The option available on a remote access connection in which you can configure the dial-up connection so that a user can dial in from only one specific phone number.

Extensible Authentication Protocol (EAP) — An authentication protocol available in Windows 2000. EAP is different from the other options in that it is really a standard for designing other authentication processes, such as certificate-based authentication.

Internet Authentication Service (IAS) — Microsoft's implementation of the Remote Authentication Dial-in User Service (RADIUS).

Internet Connection Sharing (ICS) — An option that can be configured on a remote access connection that makes it possible for multiple users in a small office to share a single connection to the Internet or another network.

Layer Two Tunneling Protocol (L2TP) — One of the virtual private network protocols supported by Windows 2000.

Microsoft Challenge Handshake Authentication Protocol (MS-CHAP) — One of the authentication protocols available with Windows 2000 remote access. It is based on a challenge and response handshake for authentication. MS-CHAP is quite secure because the password is never sent on the network in an unencrypted form.

MS-CHAP version 2 — Similar to MS-CHAP, but with enhanced security options.

Network Access Server (NAS) — The remote access server that clients connect to in a RADIUS implementation. The NAS acts as the RADIUS client in the authentication process by forwarding all authentication requests to the RADIUS Server.

remote access policies — Policies that can be configured on a Windows 2000 RRAS server that enable enhanced options for managing remote access user connections.

remote access service (RAS) — A service that is used to give remote access clients access to a network by using either a dial-up or VPN connection.

Remote Authentication Dial-in User Service (RADIUS) — An open standard that is widely used for remote access authentication in large, distributed environments such as ISPs. The purpose of RADIUS is to centralize the management of remote access users and policies in an environment in which the remote access servers might be widely distributed.

Routing and Remote Access Service (RRAS) — The Windows 2000 component that provides the remote access service, as well as routing functionality.

social engineering attack — An attack aimed at the users of a technology rather than at the technology itself. The most common way to launch a social engineering attack is to contact the user and present false credentials.

Virtual Private Network (VPN) — Used to ensure that data communication on a public network, such as the Internet, is secure. This is implemented by having the VPN create an encrypted tunnel through the Internet.

Review Questions

1. Under what circumstances should you implement a remote account lockout policy rather than just a domain level lockout policy?
   a. when you want remote users to be locked out after the same number of logon attempts as local users
   b. when you want remote users to be locked out after fewer logon attempts than local users
   c. when you want remote users to be locked out after more logon attempts than local users
   d. when you want to lock out all remote access users

2. What should you do to protect your network against a social engineering attacks?
   a. Block all incoming ports on the RAS computer.
   b. Set the remote access lockout policy at a low number.
   c. Use remote access policies to restrict access
   d. Educate the end users in your company.
3. Which type of server is the most likely to receive an attack from outside your network?
   a. application server
   b. domain controller
   c. file and print server
   d. RAS server
   e. kiosk
4. Which type of remote access would a home-based user with a high speed Internet connection use to connect to the office?
   a. client-to-server dial-up
   b. router-to-router dial-up
   c. client-to-server VPN
   d. router-to-router VPN
5. Which type of remote access would you configure for a company with two locations that need to transfer small amounts of data between offices?
   a. client-to-server dial-up
   b. router-to-router dial-up
   c. client-to-server VPN
   d. router-to-router VPN
6. You have configured a remote access policy that specifies that the members of the Managers group can dial up your RRAS server between 6:00 PM and 12:00 AM The only other remote access policy is the default policy. However, you notice that a user who is not in the Managers group can connect to the RRAS server. Where is the configuration error?
   a. The configuration settings on the remote access policy are wrong.
   b. The domain is not configured for native mode.
   c. The user is granted remote access permission in Active Directory.
   d. The RRAS server is configured incorrectly.

7. Which security strategy would not be appropriate for a traveling sales force using dial-up remote access?
   a. Set a password policy to require strong passwords.
   b. Set laptop screensavers to turn on and lock the workstation after two minutes.
   c. Enable remote access only for the sales persons user accounts.
   d. Enable callback to a preset phone number.
8. Which security strategies would be appropriate for users accessing company servers over a VPN from their home computers?
   a. Set a password policy to require strong passwords.
   b. Set a security policy that forces screensavers to turn on and lock the workstation after two minutes.
   c. Enable remote access only for the sales persons' user accounts.
   d. Enable callback to a preset phone number.
9. A Windows 2000 RAS server is capable of querying a single RADIUS server as an authentication provider. True or false?
10. What type of authentication would be required for RAS if all Windows 95 and later operating systems need to be supported?
    a. EAP
    b. MS-CHAP version 2
    c. MS-CHAP
    d. CHAP
    e. SPAP
    f. PAP
11. What type of authentication should be chosen for RAS if all clients are Windows 2000 Professional and smart cards are not used?
    a. EAP
    b. MS-CHAP version 2
    c. MS-CHAP
    d. CHAP
    e. SPAP
    f. PAP
12. Remote access policies are useful for managing remote access because:
    a. You can grant remote access permission to groups rather than individuals only.
    b. You can control what time of day users can connect to the RAS server.
    c. You can configure multilink settings.
    d. all of the above

13. What type of authentication should be chosen for RAS if some of the clients are not Windows-based and you require the highest level of security?
    a. EAP
    b. MS-CHAP version 2
    c. MS-CHAP
    d. CHAP
    e. SPAP
    f. PAP
14. You would like to configure your server to be part of a RADIUS installation hosted by your Internet service provider. Which service must be installed on a Windows 2000 server in your office?
    a. RADIUS
    b. IAS
    c. RAS
    d. Routing and Remote Access
15. You have configured RADIUS with your Internet service provider so that remote users can use the same usernames and passwords when they connect to the ISP as they do when they are at their desks. Where is the remote access server located?
    a. at the client computer at home
    b. at the client computer at the office
    c. at a Windows 2000 server at the office
    d. at the ISP
16. You decide that you would like to implement remote access policies on your RRAS server. However when you try to enable a user account to use remote access policies, the option is not available. What do you need to do?
    a. Use a Windows 2000 RRAS server.
    b. Switch the domain to native mode.
    c. Enable remote access policies for the domain controller.
    d. Install the latest service pack on the RRAS server.
17. Your company has 6 RRAS servers located in three different offices. You would like to be able to centrally manage all the remote access policies for all servers. How can you do this?
    a. Install and configure an IAS server.
    b. Move all the RRAS servers into one OU and configure the remote access policy for the OU.
    c. Configure a domain-level group policy.
    d. You can't do this.

HANDS-ON PROJECTS

Project 8-1

In this hands-on project, you will configure Routing and Remote Access as a VPN server.

To configure a VPN server:

1. Log on to your Windows 2000 computer as an administrator.
2. Click **Start**, point to **Programs**, point to **Administrative Tools**, and click **Routing and Remote Access**.
3. Click your server. Note the red down arrow indicating that the RAS is not running.
4. Right-click your server and click **Configure and Enable Routing and Remote Access**.
5. Click **Next**.
6. Select **Virtual private network (VPN) server** and click **Next**.
7. Keep the default selection of all protocols and click **Next**.
8. Select your **External** network card from the list and click **Next**.
9. Choose **From a specified range of addresses** and click **Next**.
10. Click **New**.
11. Enter a start IP address of **131.107.1.20**.
12. Enter a stop IP address of **131.107.1.30**.
13. Click **OK** and click **Next.**
14. Keep the default selection of not being a RADIUS server and click **Next**.
15. Click **Finish**. If you receive a message about relaying DHCP, then click **OK** to clear the message.
16. Expand the server in the left pane, and click **Ports**. What are the two types of VPN ports that can be connected?
17. Close all windows and log off.

Project 8-2

In this hands-on project, you will create a user account with permissions to connect to your server using a VPN connection.

To create a new user account:

1. Log on to your Windows 2000 computer as an administrator.
2. Click **Start**, point to **Programs**, point to **Administrative Tools**, and click **Active Directory Users and Computers**.
3. Right-click **Users**, point to **New**, and click **User**.

4. Name the user **VPNuser1** with the same logon name. Click **Next**.
5. Set a password of **connect** and click **Next**. Click **Finish**.
6. To configure VPNuser1 Dial-in options, give **VPNuser1** permission to access the network remotely by right-clicking **VPNuser1** and clicking **Properties**.
7. Click the **Dial-in** tab.
8. Click the **Allow Access** radio button. Click **OK**.
9. Close all windows.

Project 8-3

In this hands-on project, you will test the VPN by connecting to your VPN server using the user account created in the previous project.

To create a client VPN connection:

1. Click **Start**, point to **Settings**, and click **Network and Dial-up Connections**.
2. Double-click **Make New Connection**. Specify location information if prompted, then click **OK** to close the Phone and Modem options.
3. Click **Next** to start the **Network Connection Wizard**.
4. Select **Connect to a private network through the Internet** and click **Next**.
5. Enter in the IP address of **External** network card and click **Next**.
6. Select **Only for myself** and click **Next**. Click **Finish**.
7. To connect to a VPN server, at the VPN logon prompt, enter **VPNuser1** as the User name.
8. Enter **connect** as the password.
9. Click **Connect**.
10. Click **OK** to close the Connection Complete window if it appears.
11. View the status of the VPN connection by right-clicking **Virtual Private Connection** (found on the desktop or as the connection icon in the task bar tray) and click **Status**. Note the status of the connection, as well as the amount of traffic that has gone through the connection.
12. Click **Properties**.
13. Click the **Security** tab.
14. Select **Advanced**.
15. Click **Settings**. Note that there are varying levels you can choose for data security, as well as logon security (authentication).
16. Click **Cancel** twice.
17. Disconnect the VPN connection by clicking **Disconnect**.

Project 8-4

In this hands-on project, you will configure and test a RAS policy that only allows access between 6 PM and 8 PM Monday to Friday.

To configure a RAS policy:

1. Click **Start**, point to **Programs**, point to **Administrative Tools**, and click **Routing and Remote Access**.
2. Click the plus sign next to your server container.
3. Click **Remote Access Policies**. Note that there is one policy in place by default.
4. Right-click **Remote Access Policies** and click **New Remote Access Policy**.
5. Enter **Evening Only** as the Policy friendly name and click **Next**.
6. In the Conditions dialog box, click **Add**.
7. Double-click **Day-And-Time-Restriction**.
8. Allow access from **6PM to 8PM, Monday to Friday** and click **OK**
9. Click **Next**.
10. Choose **Grant remote access permissions if the conditions match** and click **Next**. Click **Finish**.
11. Right-click the **Evening only** policy and choose **Move Up**.
12. To test the new policy, click **Start**, point to **Programs**, point to **Administrative Tools**, and click **Active Directory Users and Computers**.
13. Right-click **VPNuser1** and choose **Properties**.
14. Click the Dial-in tab, and select the **Control access through Remote Access Policy** radio button.
15. Close all windows.
16. Double-click your VPN connection.
17. Enter **VPNuser1** as the User name.
18. Enter **connect** as the password.
19. Click **Connect**. You will get an error indicating that you do not have permission to dial in. Why do you get this error?
20. Click **Cancel**.
21. To disable Routing and Remote Access, click **Start**, point to **Programs**, point to **Administrative Tools**, and click **Routing and Remote Access**.
22. Right-click your server and click **Disable Routing and Remote Access**.
23. Click **Yes**.
24. Close all windows and log off.

CASE PROJECTS

Case Project 8-1

Up to this point, remote access has not been implemented at Southdale Property Management. However the management would like you to evaluate the possibility of implementing remote access. Many of the managers would like to be able to access e-mail when out of the office. As well, the three people who spend time outside the office would like to be able to connect to the office to collect e-mail and to access files on the company file and print server. The management is not sure how many users should be given remote access and how the remote access should be set up. Many users have high-speed Internet connections at home and have never used a dial-up connection to the Internet. While the management is interested in remote access to the Internet, they are also very concerned about any security breaches. They have asked you to make a presentation at a management meeting outlining the benefits of remote access, the security issues that need to be considered, and your recommendation for implementing remote access. What would be the main points in your presentation?

Case Project 8-2

Fleetwood Credit Union management is concerned that the existing access through dial-up phone lines is not the best way to provide remote access. The users that are connecting to the RAS server complain that the connection is very slow and not entirely reliable. As well, whenever the executives travel to other cities, the cost of the long distance call to the RAS server is quite expensive. The people outside the office all require access to their e-mail, and sometimes they also need access to files on the file server, or access to a financial application on one of the servers. You would like to implement a highly secure remote access solution that addresses all of these concerns.

In addition, management has been discussing the option of having credit union members use VPN to connect to the Fleetwood Credit Union LAN in order to secure access to the new Web financial application. Management has been told this will keep user passwords secure and prevent hackers from seeing personal information as it is transferred across the Internet. They would like your opinion on the best way to provide members with secure access to this application.